

Simulation of stochastic systems via polynomial chaos expansions and convex optimization^{*†}

Lorenzo Fagiano[‡] and Mustafa Khammash[§]

Abstract – Polynomial Chaos Expansions represent a powerful tool to simulate stochastic models of dynamical systems. Yet, deriving the expansion’s coefficients for complex systems might require a significant and non-trivial manipulation of the model, or the computation of large numbers of simulation runs, rendering the approach too time consuming and impracticable for applications with more than a handful of random variables. We introduce a novel computationally tractable technique for computing the coefficients of polynomial chaos expansions. The approach exploits a regularization technique with a particular choice of weighting matrices, which allow to take into account the specific features of Polynomial Chaos expansions. The method, completely based on convex optimization, can be applied to problems with a large number of random variables and uses a modest number of Monte Carlo simulations, while avoiding model manipulations. Additional information on the stochastic process, when available, can be also incorporated in the approach by means of convex constraints. We show the effectiveness of the proposed technique in three applications in diverse fields, including the analysis of a nonlinear electric circuit, a chaotic model of organizational behavior, finally a chemical oscillator.

^{*}This manuscript is a preprint of a paper submitted to Physical Reviews E and is subject to American Physical Society copyright. If accepted, the copy of record will be available at <http://pre.aps.org>.

[†]This research has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n. P10F-GA-2009-252284 - Marie Curie project “ICIEMSET”, from the Air Force Office of Scientific Research under the MURI award n. FA9550-10-1-0143, from the National Science Foundation through Grants ECCS-0835847 and ECCS-0802008 and from the Institute for Collaborative Biotechnologies through Grant DAAD19-03-D-0004 from the US Army Research Office.

[‡]Dip. di Automatica e Informatica, Politecnico di Torino, Italy, and Dept. of Mechanical Engineering, University of California at Santa Barbara, USA. E-mail: lorenzo.fagiano@polito.it.

[§]Dept. of Mechanical Engineering, University of California at Santa Barbara, USA, and Dept. of Biosystems Science and Engineering, ETH Zurich, Switzerland. E-mail: mustafa.khammash@bsse.ethz.ch.

1 Introduction

In most science and engineering applications, there is the need to simulate mathematical models of the process under study, in the form of ordinary or partial differential equations, with the aim to figure out the time (or space) course of a variable of interest v for analysis, decision-making and control. In many cases, in order to take into account the presence of uncertainty, unknown external inputs and in general any effect that produces a mismatch between the model equations and reality, the model to be simulated is not fully deterministic: uncertainty and disturbance are often modeled as quantities with stochastic nature, named here “input random variables” and indicated with θ , hence the name stochastic models. In these cases, $v(\theta)$ is a random variable, too, and one is interested in computing its statistics. The issue of simulating complex, nonlinear stochastic models with sufficiently high accuracy and low computational effort is still a challenge in important and diverse fields, like analysis of large power grids, weather forecasts at different scales and simulation of biological systems, to name just a few. The typical approach followed to simulate a stochastic model is the well-known Monte Carlo (MC) technique, which relies on the sampling of a finite number M of values of θ , according to its distribution. With sufficiently large M , the MC approach gives good statistical estimates (e.g. first and second order moments) of the variables of interest, and also of its probability density function (pdf). However, the application of MC simulations with the system model may be too computationally demanding, particularly in those cases when the model is complex and the inherent variables have large dimensions. Polynomial Chaos Expansions (PCEs) (see e.g. [1, 2, 3, 4, 5, 6, 7, 8, 9]) provide a useful tool to significantly reduce the computational effort required to simulate a stochastic system, by conceptually replacing the mapping between θ and v , implicitly defined by the integration of the model’s differential equations, with an explicit function $\hat{v}(\theta)$, which takes the form of a truncated series of polynomials. The polynomials in the PCE are orthogonal, so that the statistical moments of $\hat{v}(\theta)$ can be computed directly from the expansion’s coefficients. Moreover, the computational effort required to evaluate a PCE is often orders of magnitude lower than the one required to simulate the system model:

therefore, it is possible to estimate the pdf of v by using a Monte Carlo approach with the PCE, instead of the system model, with significant time savings. As an example, in the first case study considered in this paper, 100,000 Monte Carlo simulations with the system model require 3,800 s, while the same number of PCE evaluations are obtained in 24 s, with the same hardware and software. PCEs have been used with good results in a number of different areas, including experimental modeling, materials, mechanics, power systems, systems biology, and control, see e.g. [10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 22]. Yet, the computation of the PCE's coefficient may not be a trivial task when the model is nonlinear and/or the dimension n of random input variables is relatively high. In these cases, the existing approaches may require a significant and non-trivial manipulation of the model, or the computation of large numbers of simulation runs, such that the advantages of the method with respect to standard MC simulations might be lost. We propose here a novel approach to derive the PCE's coefficients, based on convex optimization. Under mild assumptions, this method can be easily applied to any existing model, since it just requires the preliminary computation of a small number of sampled values of v . In this approach, a relatively large order of the polynomial chaos is initially chosen, hence a high number of terms in the expansion, and the PCE's coefficients are then computed by means of a single multi-objective optimization problem, exploiting ℓ_1 regularization techniques (see e.g. [23],[24]). We provide a new, systematic way, particularly suited to the properties of Polynomial Chaos Expansions, to choose the weighting matrices in the cost function of the optimization problem. Moreover, we show how different kinds of available information on the stochastic model, including bounds on v and on its variance, can be easily taken into account as convex constraints in the optimization. As a result, the method is able to provide an accurate description of the statistics of $v(\theta)$, with few simulation runs. We present three case studies in a broad range of different fields, to demonstrate the effectiveness of the method and its ease of implementation.

2 Problem Settings

We consider a time-invariant system in state-space form:

$$\begin{aligned} \frac{d\mathbf{x}(t)}{dt} &= \mathbf{f}(\mathbf{x}(t), \mathbf{w}(t), \mathbf{c}, t) \\ v &= v(T) = h(\mathbf{x}(T), \mathbf{w}(T), \mathbf{c}, t) \end{aligned} \quad (1)$$

where $t \in \mathbb{R}$ is the time variable, $\mathbf{x}(t) \in \mathbb{R}^{n_x}$ is the system state, $\mathbf{w}(t) \in \mathbb{R}^{n_w}$ is an unknown input, $\mathbf{c} \in \mathbb{R}^{n_c}$ is an unknown parameter vector, finally $v \in \mathbb{R}$ is a variable of interest, evaluated at a given time instant T . $\mathbf{w}(t)$ and \mathbf{c} are assumed to have stochastic nature, in a sense that

will be better detailed afterwards. Bold symbols indicate vectors of variables, e.g. $\mathbf{x} = [x_1, \dots, x_{n_x}]^T$, where T is the vector transpose operation. The aim is to derive an approximation of the first and second order moments and the pdf of v , starting from (possibly stochastic) initial conditions $\mathbf{x}(0, \mathbf{c})$, by using the model (1). The variable \mathbf{c} accounts for uncertainty both in the model equations (e.g. due to uncertain physical parameters) and in the initial state $\mathbf{x}(0, \mathbf{c})$, while $\mathbf{w}(t)$ accounts for unknown external inputs, like disturbances. The parameters \mathbf{c} and the unknown input $\mathbf{w}(t)$, $t \in [0, T]$ in (1) can be typically expressed as functions of a n -dimensional vector $\boldsymbol{\theta} \in K^n$ of independent and identically distributed (iid) random variables θ_i , with known pdf f_θ , such that $\theta_i \in K \subset \mathcal{L}^2(\Omega, \mathcal{F}, P)$, $\forall i \in \{1, \dots, n\}$. Here, (Ω, \mathcal{F}, P) is a probability space, Ω is the set of elementary events, \mathcal{F} is the σ -algebra of the events and P is the probability measure. The expectation (or first-order moment) of a generic random variable $\theta : \mathcal{A} \rightarrow \mathbb{R}$ is denoted as $E[\theta] \doteq \int_\Omega \theta(\omega) dP(\omega) = \int_{\mathcal{A}} \theta dF_\theta$, where $F_\theta(k) \doteq P\{\theta < k\}$ is the probability distribution function of θ over \mathcal{A} . $\mathcal{L}^2(\Omega, \mathcal{F}, P)$ is the Hilbert space of all random variables θ whose L_2 -norm, $\|\theta\|_2 \doteq E[|\theta|^2]^{1/2}$, is finite, where $|\cdot|$ denotes the absolute value. K is a subspace of $\mathcal{L}^2(\Omega, \mathcal{F}, P)$ that contains only centered random variables (i.e. $\forall \theta \in K, E[\theta] = 0$). Finally, the pdf of θ is given by $f_\theta(k) = dF_\theta/dk$, and the variance (or second-order moment) of θ is indicated as $\text{Var}(\theta) \doteq E[(\theta - E[\theta])^2] = \sigma_\theta^2$, where σ_θ is the standard deviation of θ .

In many cases of practical relevance, the time-invariant parameters \mathbf{c} are naturally iid variables (e.g. when \mathbf{c} stands for some variation of a physical parameter of the system, that is uncertain due to production variability). If the variables c_i , $i \in \{1, n_c\}$ have different probability distributions, it is possible to map a standard (i.e. with zero mean and unit variance) distributed Gaussian random variable, indicated as $\mathcal{N}(0, 1)$, to a random variable with distribution function F_c by the transformation $F_c^{-1}(\text{erf}(\mathcal{N}(0, 1)))$, where erf is the Gaussian distribution function, see e.g. [25]. As regards the input $\mathbf{w}(t)$, $t \in [0, T]$, this can typically be modeled as a stochastic process or random field $\mathbf{w}(t, \boldsymbol{\theta}) : \mathbb{R} \times K^n \rightarrow \mathbb{R}^{n_w}$, which can be represented as a finite series of n iid random variables multiplied by deterministic functions $\hat{\mathbf{w}}_i(t)$, $i \in \{1, n\}$, i.e. $\mathbf{w}(t, \boldsymbol{\theta}) \simeq \hat{\mathbf{w}}_0(t) + \sum_{i=1}^n \hat{\mathbf{w}}_i(t)\theta_i$ (see for example [26, 27, 28, 14]).

We assume that the solution of the dynamical equations (1) in the time interval $[0, T]$ exists and it is unique almost surely, i.e. with probability one. In the described context, the variable of interest is a random variable, $v(\boldsymbol{\theta})$, and we name the system (1) “stochastic system”. We assume that $v(\boldsymbol{\theta})$ has finite variance:

Assumption 1 (*Finiteness of variance of $v(\boldsymbol{\theta})$*)
 $v(\boldsymbol{\theta}) \in \mathcal{L}^2(\Omega, \mathcal{F}, P)$.

Assumption 1 is typically satisfied in practical applications. The problem of simulating a stochastic system may be very complex and the main technique employed so far in engineering applications is the well-known MC approach, which consists in the following steps:

Algorithm 1 (*Monte Carlo simulations*)

1. extract M iid samples $\tilde{\theta}_{(r)}$, $r \in \{1, M\}$, of θ , according to its distribution;
2. for each sample, compute (or numerically simulate) the solution of (1) and the corresponding value of $v(\tilde{\theta}_{(r)})$, $\forall r \in \{1, M\}$;
3. analyze the statistics of the collected data.

Remark 1 For simplicity and without loss of generality, we consider a scalar variable v and a single value of T : multiple variables of interests v_j , $j \in \{1, n_v\}$ and time values T_i , $i \in \{1, N\}$ can be easily treated by considering each variable and each time instant separately from the others, as it is done in the case studies reported in this paper. In these cases, a single simulation of the system provides all of the corresponding samples $v_j(T_i, \tilde{\theta}_{(r)})$, $\forall j \in \{1, n_v\}$, $\forall i \in \{1, N\}$.

Although MC simulations proved to be very effective in many applications, the required computational times may be prohibitive in various cases, e.g. when a decision has to be taken in relatively little time on the basis of the simulations' outcome, or when repeated MC simulations have to be carried out to tune some input or parameter, or when the simulation has to be embedded in a numerical optimization procedure (e.g. for optimal design or control of stochastic systems). Polynomial Chaos Expansion techniques (see e.g. [3]) are able to significantly reduce the computational effort required by standard MC approaches, by replacing the simulation of a (possibly very complex) dynamical system with the evaluation of a static function $\hat{v}(\theta) \approx v(\theta)$. The main features of PCEs are recalled in the next Section.

3 Polynomial Chaos Expansions

PCEs were first introduced by Wiener [1], who considered Gaussian random variables θ . Later on, Cameron and Martin [2] showed one of the key properties of PCEs, namely their ability to uniformly approximate any random process with finite second-order moments. The polynomial chaos is an orthogonal basis of $\mathcal{L}^2(\Omega, \mathcal{F}, P)$, hence any random variable $v(\theta) \in \mathcal{L}^2(\Omega, \mathcal{F}, P)$ has the L_2 -convergent expansion [2]:

$$v(\theta) = \sum_{k=0}^{\infty} a_k \Phi_{\alpha_k}(\theta), \quad (2)$$

Table 1: Examples of orthogonal polynomials for different kinds of probability measure

Random variable	Polynomial basis
Gaussian	Hermite
Uniform	Legendre
Gamma	Laguerre
Beta	Jacobi

where the coefficients a_k are given by $a_k = \frac{E[v(\theta)\Phi_{\alpha_k}(\theta)]}{E[\Phi_{\alpha_k}(\theta)^2]}$, and $\Phi_{\alpha_k} = \Phi_{\alpha_k}(\theta_1, \dots, \theta_n)$ is the k -th multivariate polynomial in the series, corresponding to the k -th vector of indices, or "multi-index", $\alpha_k = [\alpha_{1,k}, \dots, \alpha_{n,k}]$, $\alpha_{i,k} \in \mathbb{N}$. More specifically, for a given vector of indices α_k , we have $\Phi_{\alpha_k}(\theta) \doteq \prod_{i=1}^n \Phi_{(\alpha_{i,k})}(\theta_i)$, where $\Phi_{(\alpha_{i,k})}(\theta_i)$ is the univariate polynomial of degree $\alpha_{i,k}$, chosen according to the Askey scheme [4]. As an example, Hermite polynomials are used with Gaussian input random variables. Table 1 shows the suitable orthogonal polynomials for different kinds of input random variables. The choice of the univariate polynomials is made in order to satisfy the orthogonality property:

$$E[\Phi_{(\alpha_j)} \Phi_{(\alpha_i)}] = E[\Phi_{(\alpha_i)}^2] \delta_{ij}, \quad (3)$$

where $\delta_{ij} = 1$ if $i = j$ and 0 in any other case. The coefficients of the univariate polynomials can be usually computed via a recursive equation, starting from the terms of degree 0 and 1. As an example, Legendre polynomials, which are orthogonal w.r.t. to the uniform probability distribution, can be obtained as:

$$\begin{aligned} \Phi_{(0)}(\theta) &= 1 \\ \Phi_{(1)}(\theta) &= \theta \\ \Phi_{(\alpha+1)}(\theta) &= \frac{2\alpha+1}{\alpha+1} \theta \Phi_{(\alpha)}(\theta) - \frac{\alpha}{\alpha+1} \Phi_{(\alpha-1)}(\theta). \end{aligned} \quad (4)$$

We denote with $l_k \doteq \sum_{i=1}^n \alpha_{i,k}$ the sum of the indices in the multi-index α_k , and we assume that the ordering of the multivariate polynomials Φ_{α_k} in (2), and of the related coefficients a_k , is such that $l_k \leq l_{k+1}$. For practical reasons, the series (2) is truncated by considering only the multi-indices up to a maximal total degree \bar{l} , i.e. $\forall \alpha_k : l_k \leq \bar{l}$. An example of ordering of all the multivariate polynomials corresponding to $\bar{l} = 2$, $n = 3$ is shown in Table 2. It can be clearly noted that the number of terms in the truncated series grows rapidly with n and \bar{l} . Since all the possible multi-indices α that sum up to $l \leq \bar{l}$ are considered, the total number L of terms in the truncated expansion is:

$$L = \frac{(n+\bar{l})!}{n! \bar{l}!}, \quad (5)$$

Table 2: Example of multivariate polynomials used in polynomial chaos, corresponding to $\bar{l} = 2$, $n = 3$

Order	Multi-index	Multivariate Polynomial
0	$\alpha_0 = [0, 0, 0]$	$\Phi_{\alpha_{n,0}}(\theta) = 1$
1	$\alpha_1 = [1, 0, 0]$	$\Phi_{\alpha_{n,1}}(\theta) = \Phi_{(1)}(\theta_1)$
1	$\alpha_2 = [0, 1, 0]$	$\Phi_{\alpha_{n,2}}(\theta) = \Phi_{(1)}(\theta_2)$
1	$\alpha_3 = [0, 0, 1]$	$\Phi_{\alpha_{n,3}}(\theta) = \Phi_{(1)}(\theta_3)$
2	$\alpha_4 = [2, 0, 0]$	$\Phi_{\alpha_{n,4}}(\theta) = \Phi_{(2)}(\theta_1)$
2	$\alpha_5 = [0, 2, 0]$	$\Phi_{\alpha_{n,5}}(\theta) = \Phi_{(2)}(\theta_2)$
2	$\alpha_6 = [0, 0, 2]$	$\Phi_{\alpha_{n,6}}(\theta) = \Phi_{(2)}(\theta_3)$
2	$\alpha_7 = [1, 1, 0]$	$\Phi_{\alpha_{n,7}}(\theta) = \Phi_{(1)}(\theta_1)\Phi_{(1)}(\theta_2)$
2	$\alpha_8 = [1, 0, 1]$	$\Phi_{\alpha_{n,8}}(\theta) = \Phi_{(1)}(\theta_1)\Phi_{(1)}(\theta_3)$
2	$\alpha_9 = [0, 1, 1]$	$\Phi_{\alpha_{n,9}}(\theta) = \Phi_{(1)}(\theta_2)\Phi_{(1)}(\theta_3)$

and the series takes the form:

$$\hat{v}(\theta) \doteq \sum_{k=0}^{L-1} a_k \Phi_{\alpha_k}(\theta) = \Phi(\theta) \mathbf{a}, \quad (6)$$

where $\mathbf{a} \doteq [a_0, \dots, a_{L-1}]^T$ and

$$\Phi(\theta) \doteq [\Phi_{\alpha_0}(\theta), \dots, \Phi_{\alpha_{L-1}}(\theta)] \quad (7)$$

are, respectively, the vectors of the PCE's coefficients and of the multivariate polynomials evaluated at θ . We refer to the truncated expansion $\hat{v}(\theta) \approx v(\theta)$ (6) as the PCE of the random variable $v(\theta)$. The PCE has been shown to converge exponentially in the L_2 -sense as the maximal order \bar{l} increases, see e.g. [3, 4]. By applying the orthogonality property (3), the first and second order moments of the random variable $\hat{v}(\theta)$ can be computed directly from the coefficients of its PCE, as follows:

$$E[\hat{v}(\theta)] = a_0 \quad (8)$$

$$\hat{\sigma}^2(\mathbf{a}) \doteq \text{Var}(\hat{v}(\theta)) = \sum_{k=1}^{L-1} a_k^2 E[\Phi_{\alpha_k}(\theta)^2] \quad (9)$$

where a_0 is the coefficient of the polynomial of order $l = 0$ (i.e. $\Phi_{\alpha_0} = 1$) in the PCE. As regards the practical computation of equation (9), the terms $E[\Phi_{\alpha_k}(\theta)^2]$, $\forall k \in \{1, L-1\}$ have to be computed once for all uses, and they can typically be obtained quite easily. As an example, for Legendre polynomials and uniformly distributed input random variables, note that, by considering that $f_\theta = 0.5d\theta$, the L_2 -norm squared of the multivariate polynomials is:

$$E[\Phi_{\alpha_k}(\theta)^2] = \|\Phi_{\alpha_k}(\theta)\|_2^2 = \prod_{i=1}^n \frac{1}{2\alpha_{i,k} + 1}, \quad (10)$$

$\forall k \in \{0, L-1\}$.

Similar equations can be derived for the other types of orthogonal polynomials.

Moreover, a Monte Carlo approach can be used to estimate the pdf of $\hat{v}(\theta)$ (and, hence, of $v(\theta)$) once the coefficients of its PCE are known, by simply evaluating the PCE $\hat{v}(\theta)$,

instead of simulating the model (1) at step 2) of Algorithm 1. The computational time required to evaluate the PCE is often orders of magnitude smaller than the one required to integrate numerically the model (1), hence the advantage of using polynomial chaos.

Clearly, one of the crucial points in the use of PCEs for the simulation of stochastic systems is the computation of the expansion's coefficients, \mathbf{a} . In the literature, this task is carried out essentially in two different ways. A first method (see e.g. [4]) relies on a Galerkin projection to obtain an augmented set of deterministic differential equations, which can be solved to compute the PCE coefficients. While this method is quite attractive from a theoretical point of view, it might be affected by some practical issues, since for complex nonlinear models it may be difficult and too time-consuming to derive the augmented set of differential equations, and the number of such equations may be too large to obtain an efficient numerical solution with standard ODE solvers.

A second approach is known as Probabilistic Collocation Method (PCM, see e.g. [6, 21]), and it basically consists in the estimation of the coefficients from a finite number of data, i.e. of ν exact values of v , named "collocation points", corresponding to ν values of the input random variables, $\tilde{\theta}_{(r)}$, $r \in \{1, \nu\}$. Here, we consider a PCM-like approach for the computation of PCEs, since it appears to be more viable for the analysis of large-scale, complex stochastic dynamical systems, and we propose a new method to estimate the coefficients. The method and its features are described in the next section. One of the main advantages of probabilistic collocation is that no modification to the original model (1) is required, but just a series of preliminary simulations to collect the data to be used in the coefficients' computation; one of its main disadvantages is that the number of collocation points can be very high, for problems with relatively high stochastic dimensions (i.e. high values of n) and strong nonlinearities. Here, we will show, through a series of case studies, that our method yields very good results even with a very low number of collocation points.

4 Computation of Polynomial Chaos Expansions via convex optimization

Given the maximal order \bar{l} of the PCE and the corresponding number of terms L (5), we propose the following algorithm to estimate the PCE's coefficients \mathbf{a} :

Algorithm 2 (*PCE computation via convex optimization*)

1. sample a finite number ν of independent values of the vector of input random variables $\tilde{\boldsymbol{\theta}}_{(r)}$, $r \in \{1, \nu\}$, according to its distribution;
2. carry out ν simulations of the system (1), each one corresponding to one of the extracted samples $\tilde{\boldsymbol{\theta}}_{(r)}$;
3. collect the obtained values of the variables of interest in the vector $\tilde{\mathbf{v}} \doteq [v(\tilde{\boldsymbol{\theta}}_{(1)}), \dots, v(\tilde{\boldsymbol{\theta}}_{(\nu)})]^T$;
4. select the maximal order \bar{l} for the PCE $\hat{v}(\boldsymbol{\theta})$ (6) and compute the matrix

$$\tilde{\Phi} \doteq \begin{bmatrix} \Phi(\tilde{\boldsymbol{\theta}}_{(1)}) \\ \vdots \\ \Phi(\tilde{\boldsymbol{\theta}}_{(\nu)}) \end{bmatrix},$$

where the vectors $\Phi(\tilde{\boldsymbol{\theta}}_{(r)})$, $r \in \{1, \nu\}$, are computed according to (7);

5. solve the following convex optimization problem to compute the PCE's coefficients:

$$\min_{\mathbf{a} \in \mathbb{R}^L} \|\mathbf{W}\mathbf{a}\|_1 + \beta \|\tilde{\Lambda}(\tilde{\mathbf{v}} - \tilde{\Phi}\mathbf{a})\|_2 \quad (11a)$$

subject to

$$\text{convex constraints} \quad (11b)$$

In (11a), for a generic vector $\mathbf{y} \in \mathbb{R}^m$ the ℓ_1 and ℓ_2 vector norms are defined as $\|\mathbf{y}\|_1 \doteq \sum_{k=1}^m |y_k|$ and $\|\mathbf{y}\|_2 =$

$\left(\sum_{k=1}^m y_k^2\right)^{1/2}$, respectively. The diagonal weighting matrix \mathbf{W} is defined as $\mathbf{W} \doteq \text{diag}(w(l_k)) \in \mathbb{R}^{L \times L}$, where l_k is the order of the multi-index $\boldsymbol{\alpha}_k$, and $w(l_k)$, $k \in \{0, L-1\}$ is a sequence of scalar weights with the following properties:

$$\begin{aligned} w(l_k) &> 0, \quad \forall k \in \{0, L-1\}; \\ w(l_k) &> w(l_j) &\iff l_k > l_j, \quad \forall k, j \in \{0, L-1\}; \\ \max_k w(l_k) &= 1. \end{aligned} \quad (12)$$

Moreover, $\beta > 0$ is a scalar weight. $w(\cdot)$ and β are parameters chosen by the user. Finally, the diagonal matrix $\tilde{\Lambda} \doteq \text{diag}(\tilde{\lambda})$ contains the values $\tilde{\lambda} \doteq [f_{\boldsymbol{\theta}}(\tilde{\boldsymbol{\theta}}_{(1)}), \dots, f_{\boldsymbol{\theta}}(\tilde{\boldsymbol{\theta}}_{(\nu)})]^T$ of the pdf $f_{\boldsymbol{\theta}}$, evaluated at the considered samples $\tilde{\boldsymbol{\theta}}_{(r)}$, $r \in \{1, \nu\}$.

We denote with \mathbf{a}^* a global minimizer of (11) and with $\hat{v}^*(\boldsymbol{\theta})$ the related PCE. The convex constraints (11b) are optional and they will be better specified later on; we now consider the unconstrained problem (11a) and discuss its features.

Weighted ℓ_1 -norm regularization of the coefficients

The convex cost function in (11a) is the weighted sum of two terms. The first one, $\|\mathbf{W}\mathbf{a}\|_1$, is a weighted ℓ_1 -norm of the PCE coefficients, in which the weights increase monotonically with the order of the related multivariate polynomials (see (12)). ℓ_1 -norm regularization is a well-established technique in function approximation and regression analysis [23],[24], and it is a convex relaxation of the problem of computing an approximation which is *sparsest*, i.e. with minimal number of non-zero terms or, equivalently, with minimal ℓ_0 quasi-norm, see e.g. [29]. However, the use of the weighting matrix \mathbf{W} is novel and pertains to the particular properties of polynomial chaos expansions. In practice, minimization of the term $\|\mathbf{W}\mathbf{a}\|_1$ yields an estimated coefficients' vector in which the terms related to higher-order multivariate polynomials have smaller absolute value. The reason for including this term in the cost function (11a) is twofold: on the one hand, it accounts for the fact that, due to the exponential convergence property of polynomial chaos, the absolute values of the PCE coefficients should decrease as the order of the corresponding polynomials in the expansion increases; on the other hand, it avoids over-determination of the fitting problem, when the number ν of data is lower than the number L of coefficients. Indeed, the ℓ_1 -norm regularization allows one to select an initial "high" maximal order \bar{l} , and then to rely on the convex optimization procedure to correctly "pick" the terms that have higher relevance, even if the number ν of sampled data points is much lower than the number of coefficients. We note that an approach with a similar goal, i.e. to obtain PCEs with low number of nonzero coefficients, has been also proposed in [30, 9], by means of an iterative algorithm, in which the maximal order \bar{l} and the number ν of data are gradually increased until a satisfactory accuracy is reached. Here, the coefficients are computed in "one shot", through the solution of a single convex optimization problem, without any iterative algorithm, and the employed number of data can be very low.

Weighted ℓ_2 -norm fitting of the data

The second term in the cost function (11a) is the empirical L_2 -norm of the error between the PCE and the computed data, i.e. $\|\tilde{\Lambda}(\tilde{\mathbf{v}} - \tilde{\Phi}\mathbf{a})\|_2 \approx E[|v - \hat{v}|^2]^{1/2} = \|v - \hat{v}\|_2$, and it accounts for the convergence property, in the L_2 -sense, of the PCE to the random variable v . The scalar weight β can be used to achieve a tradeoff between the accuracy of the PCE, with respect to the collected data, and its complexity, in terms of weighted ℓ_1 -norm. In practical applications,

with a “high” value of β the algorithm 2 yields good results and it is quite robust with respect to different choices of weights $w(l)$, provided that the properties (12) are satisfied.

If no constraints are included, problem (11) can be cast as a quadratic program, and a global solution can be efficiently computed also with thousands of coefficients [31]. We now present two kinds of convex constraints (11b), which can be used to take into account specific additional information on the random variable v , at the cost of a possibly higher computational time. These constraints are not meant to be exhaustive of the possibilities that the convex optimization approach can open, when combined with polynomial chaos.

Explicit maximal variance constraint

Since the variance $\hat{\sigma}^2$ of the expansion is a quadratic function of its coefficients (see eq. (9)), an upper bound $\bar{\sigma}^2 \geq 0$ on the variance can be explicitly enforced by the convex constraint:

$$\hat{\sigma}^2(\mathbf{a}) \leq \bar{\sigma}^2. \quad (13)$$

This constraint is always feasible, since the value $\mathbf{a} = 0$ satisfies it.

Explicit bounds on the PCE

If some convex bounds on v are known, e.g. positiveness, these can be easily included in the problem as follows. Assume, without loss of generality, that the bounds can be expressed as $g(v) \leq 0$, where $g : \mathbb{R} \rightarrow \mathbb{R}$ is a convex function (multiple convex bounds can be reduced to this form by taking the maximum among all of them). At first, a finite number μ of further iid samples $\tilde{\theta}_{(r)}$ of θ has to be computed, together with the corresponding vectors $\Phi(\tilde{\theta}_{(r)})$, $r \in \{1, \mu\}$ (7). Then, the following μ convex constraints can be included in (11b):

$$g(\Phi(\tilde{\theta}_{(r)})\mathbf{a}) \leq 0, \forall r \in \{1, \mu\}. \quad (14)$$

Indeed, if the computed minimizer \mathbf{a}^* of (11) satisfies the sampled constraints (14), it is not guaranteed that the inequality $f(\hat{v}^*(\theta)) \leq 0$ is satisfied *with probability one*; however, some probabilistic results have been established in the context of random convex programming (see [32, 33]), and these can be used to tune the number of constraints μ . As a final remark, we note that, if the bounds on v are such that $g(0) \leq 0$, then the convex optimization problem (11) is feasible with probability one in the presence of the μ constraints (14), since these are always satisfied by the value $\mathbf{a} = 0$.

5 Application examples

In this Section, we present the results obtained by applying the proposed approach in three different fields. In particular, the first example is related to the simulation of an electric




Figure 1: RLC circuit with stochastic parametric uncertainty and stochastic input. Layout of the considered electric circuit.

circuit, with both parametric uncertainty and a stochastic input. The system has weak nonlinearities, it evolves in continuous time, it has two continuous state variables and 13 input random variables. The second example is concerned with a model for organization innovation [34]. Such a model is nonlinear, it has seven positive, continuous states and it evolves in discrete time. The number of input random variables is 12. Finally, the third example is in the field of systems biology and presents the evaluation of the effects of extrinsic noise in the simulation of a chemical oscillator. This last model is simulated through the stochastic simulation algorithm (SSA) method [35], it evolves in continuous time, it has 9 positive, discrete states and 16 input random variables. All together, these examples show how the convex optimization method can be applied in a straightforward way to problems in a broad range of fields and with significant complexity, in terms of number of input random variables, nonlinearities, and constraints on the variables of interest.

5.1 RLC circuit with stochastic parametric uncertainty and stochastic input

Consider the electric circuit depicted in Fig. 1. The system equations are

$$\begin{aligned} \dot{i}_L(t) &= -\frac{1}{L}v_C(t) - \frac{R}{L}i_L(t) + \frac{1}{L}u(t) \\ \dot{v}_C(t) &= \frac{1}{C}(i_L(t) - i_D(t)), \end{aligned} \quad (15)$$

The resistance R is assumed to be a random variable $R = R_0(1 + 0.3\theta_1)$, where $R_0 = 3.5 \Omega$ and θ_1 is a random variable with uniform distribution over $[-1, 1]$. The inductance L and the capacitance C are nonlinear functions of the current $i_L(t)$ and voltage $v_C(t)$, respectively:

$$\begin{aligned} L(i_L(t)) &= 0.5 \bar{L} (1 + \exp(a_1 i_L(t)^2)) \\ C(v_C(t)) &= 0.5 \bar{C} (1 + \exp(a_2 v_C(t)^2)), \end{aligned} \quad (16)$$

where $a_1 = -0.5 \cdot 10^8$, $a_2 = -0.5 \cdot 10^6$. As an example, the function $C(v_C)$ is depicted in Fig. 2. Moreover, the maximal values \bar{L} and \bar{C} , achieved when $v_C(t) = i_L(t) = 0$, are equal to $\bar{L} = L_0(1 + 0.2\theta_2)$, $\bar{C} = C_0(1 + 0.2\theta_3)$, where




Figure 2: RLC circuit with stochastic parametric uncertainty and stochastic input. Nonlinear characteristic of capacitance $C(v_C)$ (H).

θ_2, θ_3 are also random variables with uniform distribution over $[-1, 1]$. $\theta_1, \theta_2, \theta_3$ are assumed to be independent. The circuit is connected to a device D , which may act both as load and as generator, by applying a current $i_D(t)$. In particular, $i_D(t)$ is assumed to be a stochastic input of the form:

$$i_D(t) = a_4 \sin\left(\frac{2\pi}{a_5} t\right) + a_3 i_{D,\text{rand}}(t),$$

where $a_3 = 1 \cdot 10^{-2}$, $a_4 = 5 \cdot 10^{-3}$ A, $a_5 = 1 \cdot 10^{-2}$ s. The term $\sin\left(\frac{2\pi}{a_5} t\right)$ is a known sinusoidal component, while $i_{D,\text{rand}}(t)$ is a random process with mean $\bar{i}_D = 0$ and exponential covariance function $C_D(t_1, t_2)$:

$$C_D(t_1, t_2) = \sigma_D^2 \exp^{-\mu_D |t_1 - t_2|}, \quad (17)$$

with $\sigma_D = 1$ and $\mu_D = 50$. In order to model the random process $i_{D,\text{rand}}(t)$ as a function of a finite number of random input variables, we employ the Karhunen-Loeve (KL) expansion (see e.g. [3]) with 10 independent random variables $\theta_{D,1}, \dots, \theta_{D,10}$, uniformly distributed in the interval $[-1, 1]$:

$$i_{D,\text{rand}}(t) \simeq \bar{i}_D + \sum_{i=1}^{10} \left(\frac{\sqrt{\lambda_{D,i}}}{\sigma_{\theta_D}} g_{D,i}(t) \theta_{D,i} \right). \quad (18)$$

In (18), $\sigma_{\theta_D} = 1/\sqrt{3}$ is the standard deviation of the independent random variables $\theta_{D,i}$, and $\lambda_{D,i}, g_{D,i}(t), i \in \{1, 10\}$ are, for a given maximal time range $[-T, T]$, the first ten eigenvalues and eigenfunctions of the integral equation:

$$\int_{-T}^T C_D(t_1, t_2) g_{D,i}(t_2) dt_2 = \lambda_{D,i} g_{D,i}(t_1). \quad (19)$$




Figure 3: RLC circuit with stochastic parametric uncertainty and stochastic input. Comparison between the exact covariance function $C_D(t_1, t_2)$ of the stochastic input (solid line) and the covariance function $\tilde{C}_D(t_1, t_2)$ obtained with the Karhunen-Loeve expansion (dashed line).

In the case of the exponential covariance function (17), the eigenvalues $\lambda_{D,i}$ are computed as (see [3]):

$$\lambda_{D,i} = \frac{2\sigma_D^2 \mu_D}{\omega_{D,i}^2 + \mu_D^2},$$

where $\omega_{D,i}$ are the solutions to the following transcendental equations:

$$\begin{aligned} \mu_D - \omega_{D,i} \tan(T \omega_{D,i}) &= 0, \quad i \text{ odd} \\ \omega_{D,i} + \mu_D \tan(T \omega_{D,i}) &= 0, \quad i \text{ even.} \end{aligned}$$

The corresponding eigenfunctions are:

$$\begin{aligned} g_{D,i}(t) &= \frac{\cos(\omega_{D,i} t)}{\sqrt{T + \frac{\sin(2T \omega_{D,i} t)}{2\omega_{D,i}}}}, \quad i \text{ odd} \\ g_{D,i}(t) &= \frac{\sin(\omega_{D,i} t)}{\sqrt{T - \frac{\sin(2T \omega_{D,i} t)}{2\omega_{D,i}}}}, \quad i \text{ even} \end{aligned}$$

We are interested in simulating the system subject to the constant input $u(t) = 1 \cdot 10^{-2}$ V, starting from the steady state conditions $i_L(t_0) = 0$ A, $v_L(t_0) = 1 \cdot 10^{-2}$ V, from time $t_0 = 0$ up to time $t_N = 0.02$ s. Thus, we choose $T = 0.02$ s to compute the KL expansion. Fig. 3 shows a comparison between the exact covariance function $C_D(t_1, t_2)$ and the covariance function $\tilde{C}_D(t_1, t_2)$ obtained with the KL expansion, which is computed as $\tilde{C}_D(t_1, t_2) = \sum_{i=1}^{10} (\lambda_{D,i} g_{D,i}(t_1) g_{D,i}(t_2))$, while Fig. 4 shows, as an example, 10 different realizations of the signal $i_D(t)$. We want to analyze the statistics (first and second order moments, and pdf)




Figure 4: RLC circuit with stochastic parametric uncertainty and stochastic input. Examples of 10 different realizations of the stochastic input $i_D(t)$ computed with the Karhunen-Loeve expansion.

of the current $i_L(t_i)$ and voltage $v_C(t_i)$ for $t_i = iT_s$, $i \in \{1, 10\}$, with $T_s = 2 \cdot 10^{-3}$ s. The input random variables θ include the 3 random variables related to parametric uncertainty, $\theta_1, \theta_2, \theta_3$, plus the 10 random variables involved in the KL expansion, $\theta_{D,i}$, $i \in \{1, 10\}$. Thus, there are $n = 13$ input random variables in total, all uniformly distributed in the interval $[-1, 1]$. According to Table 1, the PCE is formulated by using Legendre polynomials, whose coefficients can be easily computed via the recursion (4). We applied the convex optimization procedure to estimate the coefficients of a different PCE for the values of $v_C(t_i, \theta)$ and $i_L(t_i, \theta)$ at all the considered time instants. In particular, we carried out $\nu = 20$ initial simulations by extracting the corresponding values of $\theta_{(r)}$, $r \in \{1, \nu\}$ according to its distribution, and we used a maximal order $\bar{l} = 2$ for the PCE. This results in $L = 105$ multivariate polynomials in the expansion. Note that the number $\nu = 20$ of considered simulations is significantly small as compared to the number of random input variables, $n = 13$, and to the number of coefficients that have to be identified, $L = 105$. We made this choice to highlight one of the main advantages of the proposed convex optimization approach, i.e. to be able to obtain quite good accuracy even with an exiguous number of data. Standard methods, like least square fitting, do not share the same feature. In order to carry out a comparison between the approach proposed here and a standard least squares technique, we used the same data $v_C(t_i, \theta_{(r)}), i_L(t_i, \theta_{(r)}), i \in \{1, 10\}, r \in \{1, \nu\}$ to identify the PCE coefficients both with our convex optimization procedure and with the following 2-norm minimization problem:

$$\min_{\mathbf{a} \in \mathbb{R}^L} \|\tilde{\mathbf{v}} - \tilde{\Phi} \mathbf{a}\|_2 \quad (20)$$

Problem (20) is a standard least-square regression and it is convex, too, however in this case it is over-determined. Moreover, it does not take into account the available information on the PCE, particularly the fact that the coefficients related to lower-order terms are likely to be more important in the expansion. The PCEs obtained with the convex optimization approach are denoted as $\hat{i}_L(t_i, \theta)$, $\hat{v}_C(t_i, \theta)$, while those obtained by means of least-squares regression are denoted as $\hat{i}_L^{LS}(t_i, \theta)$, $\hat{v}_C^{LS}(t_i, \theta)$, $i \in \{1, 10\}$.

In the convex optimization approach, the weights $w(l)$, $l \in \{0, 2\}$ have been chosen as $w(0) = 0.00025$, $w(1) = .5$, $w(2) = 1$, the scalar weight β as $\beta = 5$ and the optimization problem has been solved by using the CVX package [36] for MatLab[®].

Fig. 5 shows the courses of the estimated mean values of $v_C(t_i, \theta)$ and $i_L(t_i, \theta)$ for all of the considered time instants t_i , $i \in \{1, 10\}$, obtained with 100,000 MC simulations and with the PCEs $\hat{i}_L(t_i, \theta)$, $\hat{v}_C(t_i, \theta)$ and $\hat{i}_L^{LS}(t_i, \theta)$, $\hat{v}_C^{LS}(t_i, \theta)$, $i \in \{1, 10\}$. We recall that, for the PCE approximations, the first moment of the process is computed by simply taking, for each t_i , the coefficients of the polynomial of degree 0 in the PCEs (see (8)). It can be noted that both PCE approximations (obtained either with the convex optimization approach proposed here, or with least squares regression) give an accurate estimate of the mean of the variables of interest. However, the results concerning the variance, reported in Fig. 6, are much different: while the PCE identified with the convex optimization approach proposed here achieves very good results, as compared with the extensive MC simulations, the least squares approach shows a poor accuracy. In particular, the PCEs $\hat{i}_L^{LS}(t_i, \theta)$, $\hat{v}_C^{LS}(t_i, \theta)$, $i \in \{1, 10\}$, identified through the least-square approach, show a much lower variance with respect to the other two estimates. The variances of the PCEs have been computed with the relationship (9) (and, in particular, (10)) and they have been compared, in Fig. 6, to the empirical variance estimated by means of 100,000 MC simulations of the system model.

The poor accuracy given by the PCE obtained through least squares regression is further highlighted by the estimates of the pdf, as shown as an example in Fig. 7 for the variable $i_L(t_{10}, \theta)$: while the empirical pdf computed with the PCE $\hat{i}_L(t_{10}, \theta)$ results to be very close to the one computed with the standard MC approach, the one given by $\hat{i}_L^{LS}(t_{10}, \theta)$ is very different. Finally, as regards the computational times, the 100,000 MC simulations of the model (15) took about 3800 s on a Intel[®] Core[™] 2 Duo processor at 1.3 Ghz, with 4 GB RAM and MatLab[®] 2009, while the corresponding 100,000 evaluations of the PCE required about 24 s on the same hardware. The average time required to solve the convex optimization problem in our procedure was 0.45 s, on the same computer.

Table 3: States, equations and parameters for the stochastic model of innovative search

State variables	
Incoming ideas, II	
Internal stocks, IS	
New ideas, NI	
Organizational ideas, OI	
Testing ideas, TI	
Allocation of attention, AA	
External stocks, ES	
Model equations	
$II(t+1) = c_1 ES(t+1) AA(t+1)$	
$IS(t+1) = c_2 IS(t) + NI(t) + II(t)$	
$NI(t+1) = c_3 IS(t)$	
$OI(t+1) = c_6 NI(t+1) + c_4 IS(t+1) + c_5 II(t+1)$	
$TI(t+1) = c_7 TI(t) + OI(t+1)$	
$AA(t+1) = AA(t) + c_8 NI(t) + c_9 NI^2(t) + c_{10} TI(t) + c_{11} TI^2(t)$	
$ES(t+1) = c_{12} ES(t) - II(t)$	
Parameter mean	Parameter std. dev.
$\bar{c}_1 = 0.1375$	$\sigma_{c_1} = 0.0225$
$\bar{c}_2 = 0.2$	$\sigma_{c_2} = 0.02$
$\bar{c}_3 = 0.5$	$\sigma_{c_3} = 0.06$
$\bar{c}_4 = 0.2$	$\sigma_{c_4} = 0.02$
$\bar{c}_5 = 0.2$	$\sigma_{c_5} = 0.02$
$\bar{c}_6 = 0.5$	$\sigma_{c_6} = 0.02$
$\bar{c}_7 = 0.275$	$\sigma_{c_7} = 0.025$
$\bar{c}_8 = 0.1375$	$\sigma_{c_8} = 0.0225$
$\bar{c}_9 = -0.0150$	$\sigma_{c_9} = 0.002$
$\bar{c}_{10} = -0.0505$	$\sigma_{c_{10}} = 0.0099$
$\bar{c}_{11} = 0.00055$	$\sigma_{c_{11}} = 0.00009$
$\bar{c}_{12} = 1.0055$	$\sigma_{c_{12}} = 0.0009$

5.2 Stochastic model of innovative search

This second example is concerned with a dynamical model of how organizations pursue innovation, i.e. how they allocate attention to devise new ideas, choose part of them for potential investigation, and finally selects the concepts to be actually tested. The model has been developed by [34], it accounts 7 state variables and 12 stochastic parameters, and it has been conceived in discrete time (i.e. the state at the next observation time is a function of the state at the current observation time). An overview of the model equations and of the related parameters is given in Table 3. In particular, the organization allocates a certain quantity of attention, AA , to the search for incoming ideas, II , from an external stock of ideas and information, ES . Part of these incoming ideas is stored in the internal stocks of information, IS , from which new ideas NI are derived. The time evolution of the organizational ideas that are actually selected for possible investigation, OI , is then influenced by II , IS and NI . Part of OI feeds the testing ideas, TI , i.e. the ideas that, among the organizational ideas, the organization actually chooses to pursue. The state of the system is given by

$\mathbf{x} = [II, IS, NI, OI, TI, AA, ES]^T$ and it is a vector of non-negative variables (i.e. $\mathbf{x} \in \mathbb{R}^7$, $x_i \geq 0 \forall i \in \{1, 7\}$). The value of \mathbf{x} at the observation period $t+1$ is given by a set of six nonlinear dynamical equations and one algebraic equation, involving the state at the observation period t and 12 parameters $\mathbf{c} \in \mathbb{R}^{12}$. The parameters are supposed to be independent and distributed according to Gaussian distributions, with standard deviations σ_{c_i} and mean values \bar{c}_i , $i \in \{1, 12\}$, as shown in Table 3. Thus, in this example the input random variable $\boldsymbol{\theta}$ is a vector of 12 independent Gaussian variables with normal distribution, such that $c_i = \bar{c}_i + \sigma_{c_i} \theta_i$, $i \in \{1, 12\}$. Extensive MC simulations with this model can be obtained with quite low computational times, however this example is still interesting, from the point of view of the approach proposed in this article, since the model is nonlinear and the considered parameter variations may lead to structural changes in the stability properties of the system (chaotic behavior), from stable modes, to oscillations, to divergent modes. We are interested in computing the time course, over 30 observation periods, of the first and second order moments and of the pdf of new ideas (NI), in front of the considered variability in the model pa-




Figure 5: RLC circuit with stochastic parametric uncertainty and stochastic input. Mean values at $t = t_i, i \in \{1, 10\}$ of (a) current $i_L(t)$ and (b) voltage $v_C(t)$ obtained either with MC simulations of the model (dashed line with 'o') or with the coefficients of the term of degree 0 in the PCEs $\hat{i}_L(t_i, \theta)$, $\hat{v}_C(t_i, \theta)$ (solid line with '*') and $\hat{i}_L^{LS}(t_i, \theta)$, $\hat{v}_C^{LS}(t_i, \theta)$ (dotted line with '>').

parameters. As a matter of fact, for the observation periods $t = 1, 2$ and 3 the value of NI remains fixed at its initial condition, so only the periods $t = 4, \dots, 30$ are analyzed. Following our convex optimization procedure, we sample $\nu = 300$ values of $\theta \in [-1, 1]^{12}$ and compute the corresponding values of $NI(t, \theta)$, $t \in \{4, 30\}$, starting from the same initial condition $x(0) = [0, 0, 0, 0, 0.2, 50]^T$. We use Hermite polynomials (according to Table 1) and we consider a maximal order $\bar{l} = 3$ for the PCEs $\hat{NI}(t, \theta)$, so that each expansion has $L = 455$ terms. We note that a Galerkin projection method would lead, in this case, to 3185 discrete-time dynamical equations (i.e. seven model equations, times 455 coefficients in the PCE), while a standard least-squares regression would need at least 455 sampled

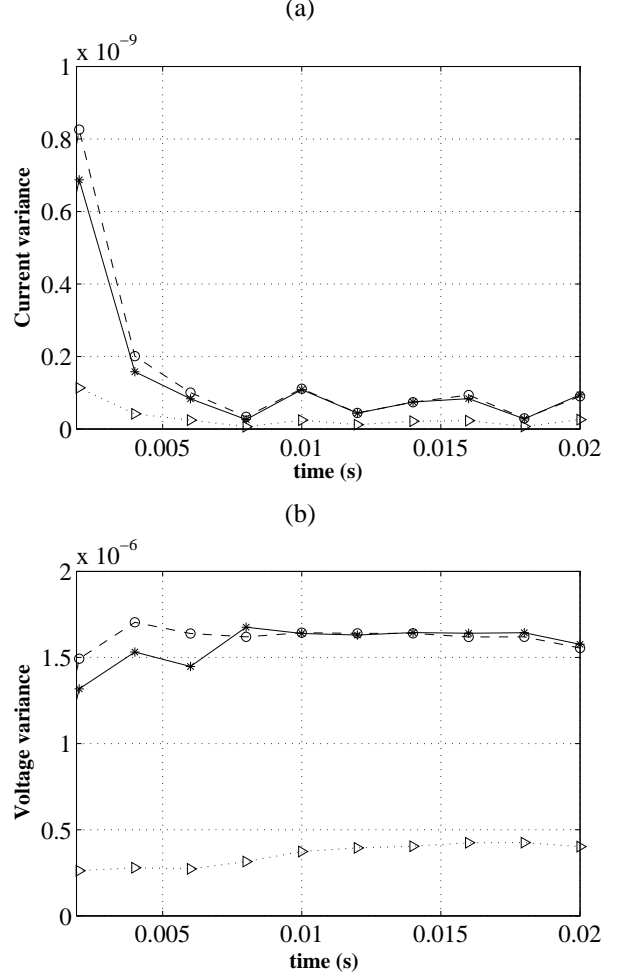


Figure 6: RLC circuit with stochastic parametric uncertainty and stochastic input. Variances at $t = t_i, i \in \{1, 10\}$ of (a) current $i_L(t)$ and (b) voltage $v_C(t)$ obtained either with MC simulations of the model (dashed line with 'o') or with the coefficients of the PCEs $\hat{i}_L(t_i, \theta)$, $\hat{v}_C(t_i, \theta)$ (solid line with '*') and $\hat{i}_L^{LS}(t_i, \theta)$, $\hat{v}_C^{LS}(t_i, \theta)$ (dotted line with '>').

data, to avoid over-determination. We select the weights $w(l)$, $l \in \{0, 3\}$ as $w(0) = 0.0001$, $w(l) = \frac{l^2}{9} \forall l \in \{1, 3\}$ and the scalar weight $\beta = 10^3$. Moreover, since variable NI is defined to be positive, we include 500 additional constraints $\hat{NI}(t, \tilde{\theta}_{(r)}) \geq 0$, $\forall r \in \{1, 500\}$, as described in the main paper. Finally, we include a quadratic constraint on the expansions' variances, by setting $\bar{\sigma}^2 = 2\bar{\sigma}^2$, where $\bar{\sigma}$ is the empiric standard deviation computed by using the 300 simulated data. We solve the convex optimization problem again with the CVX package. The courses of the first and second order moments of $NI(t, \theta)$, computed by using the PCEs' coefficients via (8)-(9), are shown in Fig. 8(a)-(b), where they are compared with the empirical moments obtained by means of 100,000 MC simulations. Fig. 9




Figure 7: RLC circuit with stochastic parametric uncertainty and stochastic input. Comparison between the empirical pdfs of variable $i_L(t_{10}, \theta)$ estimated with 100,000 MC simulations with the system model (left) and with 100,000 MC evaluations of the PCEs $\hat{i}_L(t_{10}, \theta)$ (middle) and $\hat{i}_L^{LS}(t_{10}, \theta)$ (right).

shows the time evolution of the pdf of NI , estimated either by computing 100,000 MC simulations or with the corresponding 100,000 evaluations of the PCEs. Finally, Table 4 shows, as an example, a comparison between the quartiles of the variables of interest, computed either with standard MC simulations, or with the corresponding PCE. The good matching between all of these statistics shows that, also in this case, the expansions' coefficients, computed with the proposed convex optimization method, are able to describe the stochastic process with good accuracy. This second example is concerned with a dynamical model of how organizations pursue innovation, i.e. how they allocate attention to devise new ideas, choose part of them for potential investigation, and finally selects the concepts to be actually tested. The model has been developed by [34], it accounts 7 state variables and 12 stochastic parameters, and it has been conceived in discrete time (i.e. the state at the next observation time is a function of the state at the current observation time). An overview of the model equations and of the related parameters is given in Table 3. In particular, the organization allocates a certain quantity of attention, AA , to the search for incoming ideas, II , from an external stock of ideas and information, ES . Part of these incoming ideas is stored in the internal stocks of information, IS , from which new ideas NI are derived. The time evolution of the organizational ideas that are actually selected for possible investigation, OI , is then influenced by II , IS and NI . Part of OI feeds the testing ideas, TI , i.e. the ideas that, among the organizational ideas, the organization actually chooses to pursue. The state of the system is given by $\mathbf{x} = [II, IS, NI, OI, TI, AA, ES]^T$ and it is a vector of non-negative variables (i.e. $\mathbf{x} \in \mathbb{R}^7$, $x_i \geq 0 \forall i \in \{1, 7\}$). The value of \mathbf{x} at the observation period $t + 1$ is given by a

Table 4: Stochastic model of innovative search: quartiles of the probability distribution of new ideas during 27 observation periods, estimated either with standard MC simulations, or with the PCEs computed with the convex optimization approach

Obs. period	MC simulations			Polynomial chaos		
	25%	50%	75%	25%	50%	75%
4	0.59	0.68	0.78	0.59	0.68	0.78
5	0.67	0.77	0.87	0.66	0.76	0.86
6	0.86	1.00	1.16	0.87	1.01	1.16
7	1.01	1.20	1.41	1.00	1.19	1.41
8	1.17	1.43	1.74	1.17	1.44	1.74
9	1.32	1.68	2.10	1.33	1.68	2.10
10	1.50	1.96	2.52	1.50	1.95	2.52
11	1.66	2.24	2.96	1.67	2.24	2.96
12	1.83	2.54	3.43	1.83	2.54	3.43
13	1.99	2.83	3.89	1.99	2.83	3.89
14	2.15	3.11	4.31	2.14	3.11	4.30
15	2.28	3.36	4.65	2.28	3.36	4.65
16	2.41	3.55	4.89	2.40	3.56	4.89
17	2.51	3.70	5.01	2.52	3.71	5.01
18	2.58	3.77	5.00	2.60	3.78	5.01
19	2.63	3.77	4.90	2.65	3.78	4.92
20	2.65	3.70	4.72	2.66	3.70	4.73
21	2.62	3.57	4.49	2.65	3.59	4.51
22	2.56	3.41	4.24	2.58	3.41	4.23
23	2.46	3.22	3.99	2.50	3.23	3.98
24	2.33	3.04	3.75	2.39	3.05	3.73
25	2.18	2.85	3.52	2.25	2.86	3.48
26	2.03	2.66	3.30	2.11	2.69	3.26
27	1.87	2.49	3.11	1.98	2.53	3.09
28	1.70	2.31	2.93	1.81	2.36	2.90
29	1.54	2.15	2.76	1.65	2.21	2.75
30	1.38	1.99	2.61	1.49	2.05	2.60

set of six nonlinear dynamical equations and one algebraic equation, involving the state at the observation period t and 12 parameters $\mathbf{c} \in \mathbb{R}^{12}$. The parameters are supposed to be independent and distributed according to Gaussian distributions, with standard deviations σ_{c_i} and mean values \bar{c}_i , $i \in \{1, 12\}$, as shown in Table 3. Thus, in this example the input random variable θ is a vector of 12 independent Gaussian variables with normal distribution, such that $c_i = \bar{c}_i + \sigma_{c_i}\theta_i$, $i \in \{1, 12\}$. Extensive MC simulations with this model can be obtained with quite low computational times, however this example is still interesting, from the point of view of the approach proposed in this article, since the model is nonlinear and the considered parameter variations may lead to structural changes in the stability properties of the system (chaotic behavior), from stable modes, to oscillations, to divergent modes. We are interested in computing the time course, over 30 observation periods, of the first and second order moments and of the pdf of new ideas

(NI), in front of the considered variability in the model parameters. As a matter of fact, for the observation periods $t = 1, 2$ and 3 the value of NI remains fixed at its initial condition, so only the periods $t = 4, \dots, 30$ are analyzed. Following our convex optimization procedure, we sample $\nu = 300$ values of $\theta \in [-1, 1]^{12}$ and compute the corresponding values of $NI(t, \theta)$, $t \in \{4, 30\}$, starting from the same initial condition $x(0) = [0, 0, 0, 0, 0.2, 50]^T$. We use Hermite polynomials (according to Table 1) and we consider a maximal order $\bar{l} = 3$ for the PCEs $\hat{NI}(t, \theta)$, so that each expansion has $L = 455$ terms. We note that a

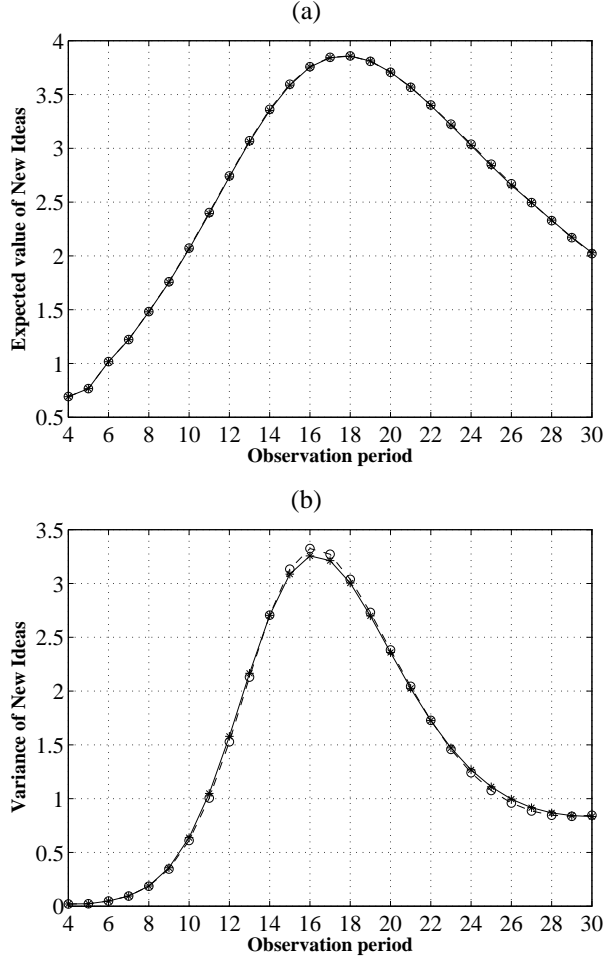


Figure 8: Stochastic model of innovative search. (a) Mean values and (b) variances of the New Ideas NI at the observation periods $t = 4, \dots, 30$, estimated either with 100,000 standard MC simulations (dashed line with 'o') or with the polynomial chaos expansions computed with the convex optimization approach (solid lines with '*').

Galerkin projection method would lead, in this case, to 3185 discrete-time dynamical equations (i.e. seven model equations, times 455 coefficients in the PCE), while a standard least-squares regression would need at least 455 sampled

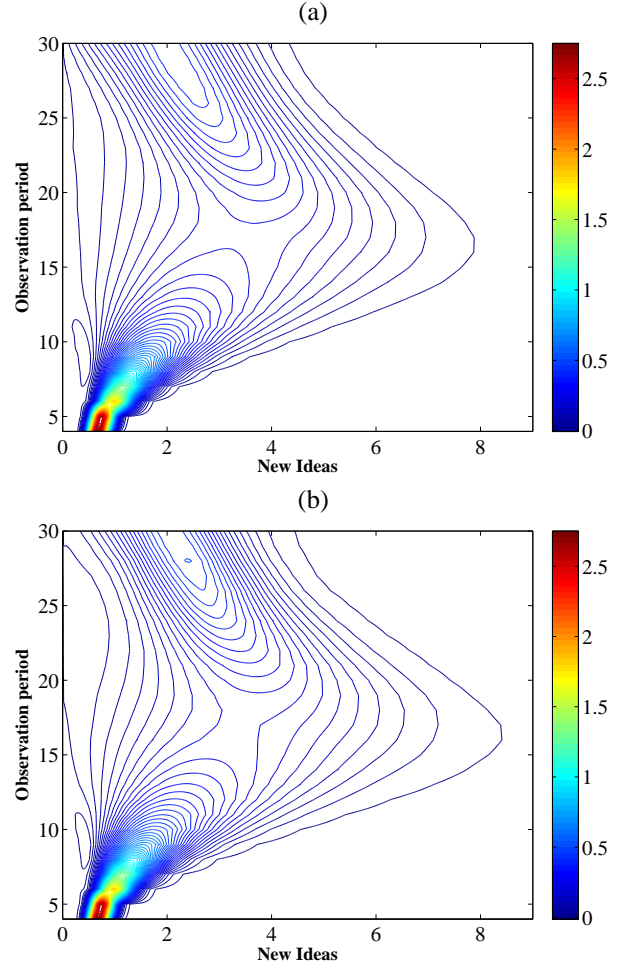


Figure 9: Stochastic model of innovative search. Level curves of the pdf of new ideas NI as a function of the observation period, estimated by means of (a) 100,000 standard MC simulations and (b) polynomial chaos expansions computed with the convex optimization approach.

data, to avoid over-determination. We select the weights $w(l)$, $l \in \{0, 3\}$ as $w(0) = 0.0001$, $w(l) = \frac{l^2}{9} \forall l \in \{1, 3\}$ and the scalar weight $\beta = 10^3$. Moreover, since variable NI is defined to be positive, we include 500 additional constraints $\hat{NI}(t, \tilde{\theta}_{(r)}) \geq 0$, $\forall r \in \{1, 500\}$, as described in the main paper. Finally, we include a quadratic constraint on the expansions' variances, by setting $\bar{\sigma}^2 = 2\tilde{\sigma}^2$, where $\tilde{\sigma}$ is the empiric standard deviation computed by using the 300 simulated data. We solve the convex optimization problem again with the CVX package. The courses of the first and second order moments of $NI(t, \theta)$, computed by using the PCEs' coefficients via (8)-(9), are shown in Fig. 8(a)-(b), where they are compared with the empirical moments obtained by means of 100,000 MC simulations. Fig. 9 shows the time evolution of the pdf of NI , estimated either by computing 100,000 MC simulations or with the corre-

sponding 100,000 evaluations of the PCEs. Finally, Table 4 shows, as an example, a comparison between the quartiles of the variables of interest, computed either with standard MC simulations, or with the corresponding PCE. The good matching between all of these statistics shows that, also in this case, the expansions’ coefficients, computed with the proposed convex optimization method, are able to describe the stochastic process with good accuracy.

5.3 Chemical oscillator

The last application example is in the field of systems biology. We consider the chemical oscillator analyzed in [37] and we simulate this system by means of Gillespie’s stochastic simulation algorithm [35] (SSA) and the common reaction path (CRP) method proposed in [38]. In this chemical process, the promoters P_R and P_A control a repressor protein R and an activator protein A , respectively. The latter is able to combine with either P_R or P_A , giving rise to the complexes P_{RA} and P_{AA} , to enhance the transcription of $mRNA_A$ or $mRNA_R$, respectively, with the consequent synthesis of new A or R proteins. The repressor protein R is able to combine with A , by forming the intermediate complex A_R , and to induce its degradation. The state of this model is given by the quantities of the involved molecules, i.e. $\mathbf{x} = [A, R, P_A, P_R, P_{AA}, P_{RA}, mRNA_A, mRNA_R, A_R]^T \in \mathbb{R}_{\geq 0}^9$, which are discrete by definition, while the model evolves in continuous time. In particular, according to the SSA-CRP simulation method, each one of the 16 reactions that may occur in this process has its own internal clock, and its own stream of random firing times, whose total number, at a given time instant, is a Poisson random variable. We indicate these streams of random firing times as $\xi = \{\xi_i \doteq \{\tau_{k_i}\}_{k_i=0}^\infty, i \in \{1, 16\}\}$, where τ_{k_i} is the time interval between two subsequent firing times for the i^{th} reaction, and $k_i \in \mathbb{N}$ is a counter giving the total number of reactions of type i that already took place. The internal clocks of the reactions evolve at different speeds, given by the propensities $a_i, i = 1, \dots, 16$ times the common time variable, t . The latter are generally nonlinear functions of the state and of 16 model parameters $c_i, i = 1, \dots, 16$. The model’s chemical reactions, the related parameters and the propensities’ equations are described in Table 5. As an example, consider the 4th reaction and assume it took place already 10 times, i.e. $k_4 = 10$. When the 4th internal clock hits its own next firing time, given by $\sum_{k_4=0}^{11} \tau_{k_4}$, the 4th reaction takes place again, the counter k_4 is augmented by one, and the system state is updated according to the corresponding chemical equation (i.e. the number of $mRNA_R$ molecules is augmented by one), as well as the values of the propensities. Then, the simulation continues with the new propensity values (i.e. the new clocks’ “speeds”). Since each re-

action, when it takes place, influences the propensities of the other reactions, the simulation must be carried out in a sequential fashion, by iteratively computing the reaction that fires next and updating the state, propensities and counters (for more details, the interested reader is referred to e.g. [38]). Typically, SSA simulations are carried out for a given initial state and fixed parameters’ values, by taking multiple random extractions of the firing times’ streams ξ (internal noise), and then some statistic of interest is analyzed. Yet, the model parameters are not fixed and known, rather they can be assumed to be random variables themselves, the so-called extrinsic noise, and it is of interest to study the sensitivity of the SSA outcome to such parameter variations. PCEs have been already applied in the context of systems biology [18], by using a projection method and then a quadrature approach to identify the PCE coefficients. As it is also recalled in [18], Gauss-Hermite quadrature is efficient up to 3-5 stochastic dimensions. In this example, we analyze the sensitivity to random perturbations in all of the 16 parameters, thus we have 16 stochastic dimensions. Each one of the 16 model parameters $c_i, i = 1, \dots, 16$, is assumed to be uniformly distributed in the interval $\pm 10\%$ centered at the corresponding mean value \bar{c}_i , as indicated in Table 5. Therefore, the input random variables are given by a vector $\boldsymbol{\theta} = \{\theta_i \in [-1, 1] : c_i = (1 + 0.1\theta_i)\bar{c}_i, i = 1, \dots, 16\}$ of 16 independent variables, each one with uniform distribution in the interval $[-1, 1]$. The variables of interest are the expected value $\bar{A}(t, \boldsymbol{\theta}) = E[A(t, \boldsymbol{\theta}, \xi)]$ and the variance $\sigma_A^2(t, \boldsymbol{\theta}) = E[(A(t, \boldsymbol{\theta}, \xi) - \bar{A}(t, \boldsymbol{\theta}, \xi))^2]$ of the number of molecules of protein A , evaluated every 5 s up to $t = 50$ s of simulation time, starting from the initial condition $\mathbf{x}(0) = [0, 177, 1, 1, 0, 0, 4, 0, 279]^T$. More specifically, we consider the empirical values of $\bar{A}(t, \boldsymbol{\theta}), \sigma_A^2(t, \boldsymbol{\theta})$, computed by averaging over 1000 SSA simulations. By using the CPR method [38], each SSA realization is associated to its own, fixed seed that generates the random streams ξ . In this way, each simulation is evaluated with different values of $\boldsymbol{\theta}$ (i.e. extrinsic noise) but always with the same random stream ξ (i.e. internal noise): therefore, extrinsic noise and internal noise are effectively decoupled, since in practice, for a fixed value of $\boldsymbol{\theta}$, the process is totally deterministic and given by the SSA simulations, each one with its own stream of random firing times. Thus, the only source of randomness lies in the model parameters $\mathbf{c} \in \mathbb{R}^{16}$. Indeed, the application of Galerkin projection methods appears to be not trivial in this case, due to the particularity of the described SSA method and to the discrete nature of the state variables. Moreover, we consider PCEs of order 3, which, in a 16-dimensional space, involve 969 terms. With the convex optimization approach, we run 100 sets of 1,000 SSA simulations, corresponding to 100 samples of the random parameter vector. This number of data is very low with respect to the dimensionality of the random parameter $\boldsymbol{\theta} \in [-1, 1]^{16}$, yet the resulting 3rd-order PCE is

Table 5: Reactions, propensities and nominal parameter values for the chemical oscillator. The stochastic model parameters are uniformly distributed in the interval $\pm 10\%$ around the nominal value

Reaction	Propensity	Nominal parameters
$P_A \xrightarrow{a_1} P_A + mRNA_A$	$a_1 = c_1 P_A$	$\bar{c}_1 = 50$
$P_{AA} \xrightarrow{a_2} P_{AA} + mRNA_A$	$a_2 = c_{15} c_1 P_{AA}$	$\bar{c}_2 = 0.01$
$P_R \xrightarrow{a_3} P_R + mRNA_R$	$a_3 = c_2 P_R$	$\bar{c}_3 = 50$
$P_{RA} \xrightarrow{a_4} P_{RA} + mRNA_R$	$a_4 = c_{16} c_2 P_{RA}$	$\bar{c}_4 = 5$
$mRNA_A \xrightarrow{a_5} mRNA_A + A$	$a_5 = c_3 mRNA_A$	$\bar{c}_5 = 20$
$mRNA_R \xrightarrow{a_6} mRNA_R + R$	$a_6 = c_4 mRNA_R$	$\bar{c}_6 = 1$
$A + R \xrightarrow{a_7} A - R$	$a_7 = c_5 A R$	$\bar{c}_7 = 50$
$P_A + A \xrightarrow{a_8} P_{AA}$	$a_8 = c_6 P_A A$	$\bar{c}_8 = 1$
$P_{AA} \xrightarrow{a_9} P_A + A$	$a_9 = c_7 P_{AA}$	$\bar{c}_9 = 100$
$P_R + A \xrightarrow{a_{10}} P_{RA}$	$a_{10} = c_8 P_R A$	$\bar{c}_{10} = 1$
$P_{RA} \xrightarrow{a_{11}} P_R + A$	$a_{11} = c_9 P_{RA}$	$\bar{c}_{11} = 0.2$
$A \xrightarrow{a_{12}} \emptyset$	$a_{12} = c_{10} A$	$\bar{c}_{12} = 10$
$R \xrightarrow{a_{13}} \emptyset$	$a_{13} = c_{11} R$	$\bar{c}_{13} = 0.5$
$mRNA_A \xrightarrow{a_{14}} \emptyset$	$a_{14} = c_{12} mRNA_A$	$\bar{c}_{14} = 1$
$mRNA_R \xrightarrow{a_{15}} \emptyset$	$a_{15} = c_{13} mRNA_R$	$\bar{c}_{15} = 10$
$A_R \xrightarrow{a_{16}} \emptyset$	$a_{16} = c_{14} A_R$	$\bar{c}_{16} = 5,000$

Table 6: Chemical oscillator: quartiles of the probability distribution of the expected number $\bar{A}(t, \theta)$ of molecules of protein A , computed every 5 s up to 50 s, estimated either with standard MC simulations, or with the PCEs computed with the convex optimization approach

Time (s)	MC simulations			Polynomial chaos		
	25%	50%	75%	25%	50%	75%
5	58	113	217	51	123	242
10	547	686	833	556	693	841
15	258	325	392	251	319	388
20	25	38	56	29	44	61
25	3	5	8	4	7	9
30	16	30	54	17	33	58
35	174	263	381	181	272	386
40	403	489	591	404	493	595
45	200	280	371	196	278	373
50	37	63	104	42	67	110

Table 7: Chemical oscillator: quartiles of the probability distribution of the variance $\sigma_A^2(t, \theta)$ of the number of molecules of protein A , computed every 5 s up to 50 s, estimated either with standard MC simulations, or with the PCEs computed with the convex optimization approach

Time (s)	MC simulations ($\times 10^3$)			Polynomial chaos ($\times 10^3$)		
	25%	50%	75%	25%	50%	75%
5	54.4	115.2	229.3	53.7	129.2	259.5
10	298.5	364.7	436.4	294.4	356.5	422.6
15	192.7	235.6	285.8	187.2	238.5	295.9
20	24.3	37.4	53.5	28.4	42.0	58.7
25	3.0	5.2	7.8	3.7	5.7	8.1
30	17.1	32.9	62.3	19.3	38.1	68.9
35	159.3	238.8	346.0	161.1	242.5	344.2
40	264.3	335.9	424.5	264.1	335.3	422.1
45	164.7	226.4	305.0	164.5	231.1	311.2
50	35.6	61.1	103.2	41.4	66.4	109.2

highly accurate with respect to the results obtained with a standard MC approach, considering 10,000 sets of 1,000 SSA simulations. In particular, the comparison between the two approaches is shown in Fig. 10-13 and in Tables 6-7. In the convex optimization problem, we chose the weights $w(l)$, $l \in \{0, 3\}$ as $w(0) = 0.0001$, $w(l) = \frac{l^3}{27} \forall l \in \{1, 3\}$, and the scalar weight $\beta = 10^3$. We solved the convex optimization problem by using the Yalmip [39] toolbox for MatLab®. Moreover, similarly to the second application example, we included 5,000 additional constraints $\hat{A}(t, \tilde{\theta}_{(r)}), \hat{\sigma}_A^2(t, \tilde{\theta}_{(r)}) \geq 0, \forall r \in \{1, 5, 000\}$, in order to take into account the positiveness of the variables of interest. As regards the computational times, the time required to compute the 100 data points used to identify the PCEs' coefficient was 67 min, the solution of the 20 convex optimization problems (2 variables of interest evaluated at 10 different time instants) took 2 hours (averagely 6 minutes for each PCE), finally the evaluation of 10,000 MC values of the resulting PCE took 8 s on an Intel® Core™ 2 Duo processor at 1.3 GHz, with 4 GB RAM and MatLab® 2009. Thus, the PCE-convex optimization method took about 3 hours in total. The time required to compute the 10000 standard MC simulations was about 6670 minutes, i.e. 4.6 days. The model equations for the SSA simulations have been programmed in Simulink®, and the computation have been carried out on a Speedgoat® real-time machine, by using Embedded Matlab® and xPC-target® tools to automatically generate the simulation code from the Simulink model. Indeed, the results of this example confirm that the proposed method, based on convex optimization, is able to compute the PCE's coefficients with good accuracy also in the presence of a relatively large number of random dimensions and model nonlinearities, with a very limited number of preliminary data. We note that, once the PCE's coefficients have been computed, 100,000 evaluations of the expansion would take, on the same Intel® Core™ 2 Duo

computer, about 80 s, while the corresponding simulations with the dedicated real-time hardware would take about 46 days.

6 Conclusions

We proposed a new method to compute polynomial chaos expansions, by means of a suitably defined convex optimization problem. The method can easily handle thousands of terms in the PCE, corresponding for example to stochastic dimensions of 15-20 with orders of 3-4. Bounds on the first and second order moments and on the values of the resulting PCE can also be explicitly included. We applied the approach to three examples in a broad range of different fields: in all cases, the derived PCEs, computed via a very low number of preliminary simulations, accurately captured the process' statistics, despite the presence of nonlinearities and high stochastic dimensions. This aspect indicates that a quite small number of sampled simulations already contains sufficient information on the process, to derive an accurate PCE approximation. This method can be straightforwardly used in a large variety of applications, since it does not require any modification to the existing model, but just a small number of simulation runs.

References

- [1] N. Wiener, "The homogeneous chaos," *American Journal of Mathematics*, vol. 60, pp. 897–936, 1938.
- [2] R. Cameron and W. Martin, "The orthogonal development of non-linear functionals in series of fourier-hermite functionals," *Annals of Mathematics*, vol. 48, pp. 385–392, 1947.
- [3] R. Ghanem and P. Spanos, *Stochastic finite elements - A spectral approach*. Springer, 1991.
- [4] D. Xiu and G. E. Karniadakis, "The wiener-askew polynomial chaos for stochastic differential equations," *SIAM Journal of Scientific Computation*, vol. 24, pp. 619–644, 2002.

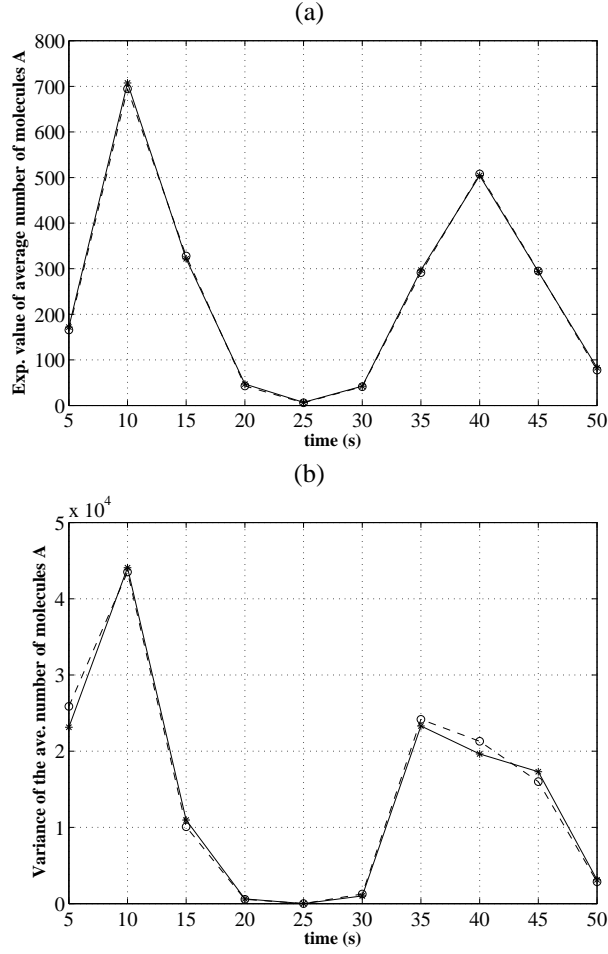


Figure 10: Chemical oscillator. (a) Mean values and (b) variances of the expected number $\bar{A}(t, \theta)$ of molecules of protein A estimated every 5 s up to 50 s of simulation, either with 10,000 standard MC simulations (dashed line with 'o') or with the polynomial chaos expansions computed with the convex optimization approach (solid lines with '*').

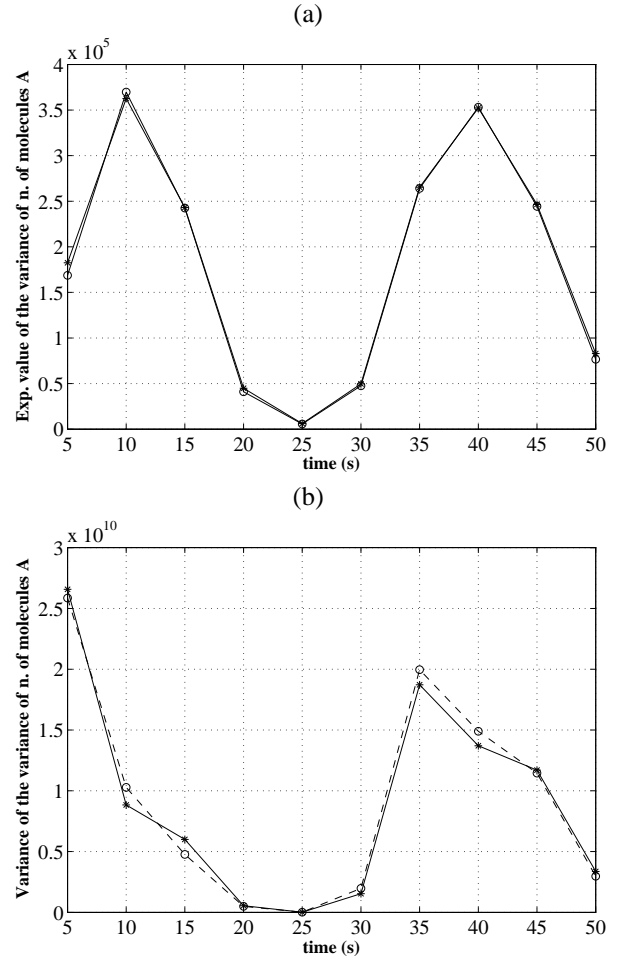


Figure 11: Chemical oscillator. (a) Mean values and (b) variances of the variance $\sigma_A^2(t, \theta)$ of molecules of protein A estimated every 5 s up to 50 s of simulation, either with 10,000 standard MC simulations (dashed line with 'o') or with the polynomial chaos expansions computed with the convex optimization approach (solid lines with '*').

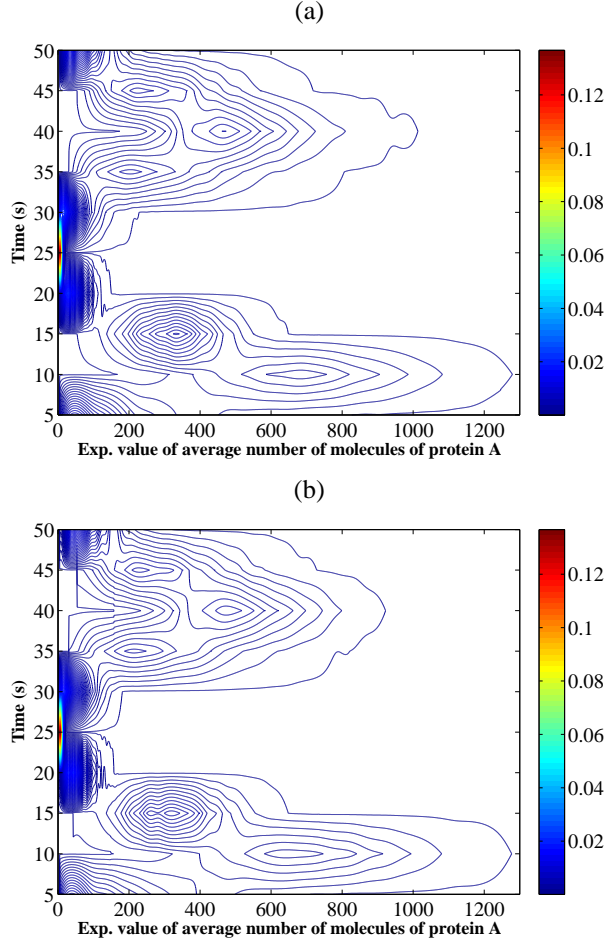


Figure 12: Chemical oscillator. Level curves of the pdf of the average value $\bar{A}(t, \theta)$ of the number of protein A molecules as a function of time, estimated by means of (a) 10,000 standard MC simulations and (b) polynomial chaos expansions computed with the convex optimization approach.

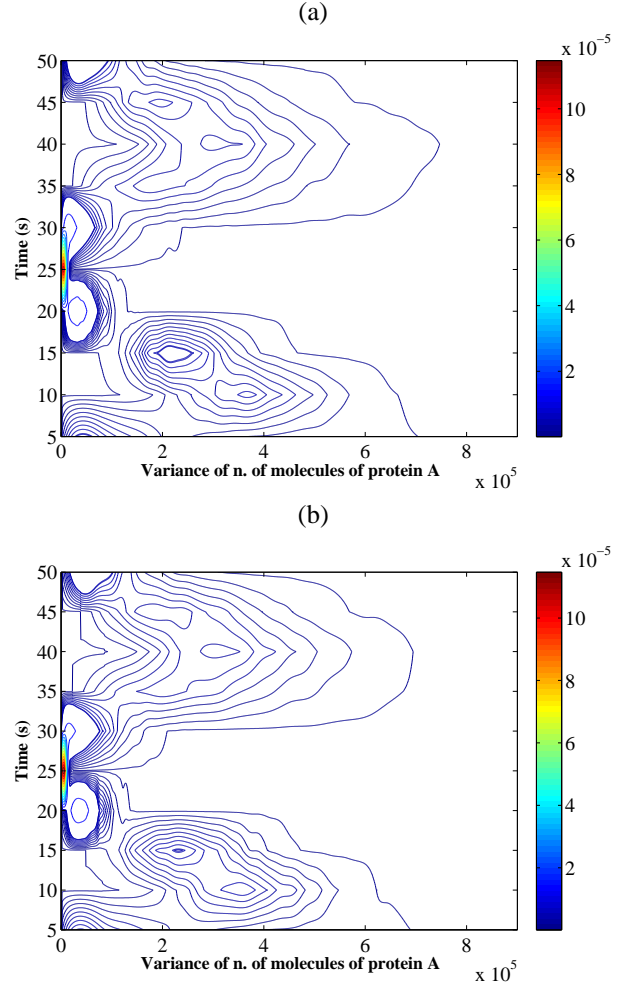


Figure 13: Chemical oscillator. Level curves of the pdf of the variance $\sigma_A^2(t, \theta)$ of the number of protein A molecules as a function of time, estimated by means of (a) 10,000 standard MC simulations and (b) polynomial chaos expansions computed with the convex optimization approach.

- [5] X. Wan and G. E. Karniadakis, "An adaptive multi-element generalized polynomial chaos method for stochastic differential equations," *Journal of Computational Physics*, vol. 209, pp. 617–642, 2005.
- [6] D. Xiu and J. S. Hesthaven, "High-order collocation methods for differential equations with random inputs," *SIAM Journal on Scientific Computing*, vol. 27, no. 3, pp. 1118–1139, 2005.
- [7] X. Wan, B. Rozovskii, and G. E. Karniadakis, "A stochastic modeling methodology based on weighted wiener chaos and malliavin calculus," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 106, no. 34, pp. 14 189–14 194, 2009.
- [8] M. Gerritsma, J. B. van der Steen, P. Vos, and G. Karniadakis, "Time-dependent generalized polynomial chaos," *Journal of Computational Physics*, vol. 229, pp. 8333–8363, 2010.
- [9] G. Blatman and B. Sudret, "Adaptive sparse polynomial chaos expansion based on least angle regression," *Journal of Computational Physics*, vol. 230, p. 2345–2367, 2011.
- [10] S. Kakutani, "Determination of the spectrum of the flow of brownian motion," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 36, no. 5, pp. 319–323, 1950.
- [11] M. Barahona and C. Poon, "Detection of nonlinear dynamics in short, noisy time series," *Nature*, vol. 381, pp. 215–217, 1996.
- [12] D. Ghiocel and R. Ghanem, "Stochastic finite element analysis of seismic soil-structure interaction," *Journal of Engineering Mechanics*, vol. 128, no. 1, pp. 66–77, 2002.
- [13] D. Xiu, D. Lucor, C. H. Su, and G. E. Karniadakis, "Stochastic modeling of flow-structure interactions using generalized polynomial chaos," *Journal of Fluids Engineering*, vol. 124, p. 51–59, 2002.
- [14] D. Lucor and G. Karniadakis, *Computational Fluid and Solid Mechanics 2003*, 2003, vol. 2, ch. Stochastic flow-structure interactions, pp. 1426–1429.
- [15] J. R. Hockenberry and B. C. Lesieutre, "Evaluation of uncertainty in dynamic simulations of power system models: the probabilistic collocation method," *IEEE Transactions on Power Systems*, vol. 19, no. 3, pp. 1483–1491, 2004.
- [16] G. Lin, C. H. Su, and G. E. Karniadakis, "The stochastic piston problem," *Proceedings of the National Academy of Sciences of the United States of America*, vol. 101, no. 45, pp. 15 840–15 845, 2004.
- [17] F. S. Hover and M. S. Triantafyllou, "Application of polynomial chaos in stability and control," *Automatica*, vol. 42, pp. 789–795, 2006.
- [18] D. Kim, B. J. Debuschere, and H. N. Najm, "Spectral methods for parametric sensitivity in stochastic dynamical systems," *Biophysical journal*, vol. 92, no. January 2007, pp. 379–393, 2007.
- [19] F. S. Hover, "Gradient dynamic optimization with legendre chaos," *Automatica*, vol. 44, pp. 135–140, 2008.
- [20] S. Das, R. Ghanem, and S. Finette, "Polynomial chaos representation of spatio-temporal random fields from experimental measurements," *Journal of Computational Physics*, vol. 228, no. 23, pp. 8726–8751, 2009.
- [21] Z. N. Nagy and R. Braatz, "Distributional uncertainty analysis using polynomial chaos expansions," in *Multi-conference on Systems and Control*, Yokohama, Japan, 2010, pp. 1103–1108.
- [22] Y. B. Peng, R. Ghanem, and J. Li, "Polynomial chaos expansions for optimal control of nonlinear random oscillator," *Journal of Sound and Vibration*, vol. 329, no. 18, pp. 3660–3678, 2010.
- [23] R. Tibshirani, "Regression shrinkage and selection via the lasso," *J. Royal. Statist. Soc. B.*, vol. 58, no. 1, pp. 267–288, apr. 1996.
- [24] D. Donoho, "Compressed sensing," *Information Theory, IEEE Transactions on*, vol. 52, no. 4, pp. 1289–1306, apr. 2006.
- [25] A. Papoulis, *Probability, Random Variables, and Stochastic Processes*, 3rd edition. Singapore: McGraw-Hill, 1991.
- [26] I. Babuška and P. Chatzipantelidis, "On solving linear elliptic stochastic partial differential equations," *Computer Methods in Applied Mechanics and Engineering*, vol. 191, pp. 4093–4122, 2002.
- [27] M. K. Deb, I. Babuška, and J. T. Oden, "Solution of stochastic partial differential equations using galerkin finite element techniques," *Computer Methods in Applied Mechanics and Engineering*, vol. 190, pp. 6359–6372, 2001.
- [28] I. Babuška, K. Liu, and R. Tempone, "Solving stochastic partial differential equations based on the experimental data," *Mathematical Models and Methods in Applied Sciences*, vol. 13, pp. 415–444, 2003.
- [29] J. Fuchs, "Recovery of exact sparse representations in the presence of bounded noise," *Information Theory, IEEE Transactions on*, vol. 51, no. 10, pp. 3601–3608, oct. 2005.
- [30] G. Blatman and B. Sudret, "An adaptive algorithm to build up sparse polynomial chaos expansions for stochastic finite element analysis," *Probabilistic Engineering Mechanics*, vol. 25, pp. 183–197, 2010.
- [31] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge university press, 2009, http://www.stanford.edu/~boyd/cvxbook/bv_cvxbook.pdf.
- [32] G. Calafiore and M. Campi, "Uncertain convex programs: Randomized solutions and confidence levels," *Mathematical Programming*, vol. 102, no. 1, pp. 25–46, 2005.
- [33] G. Calafiore, "Random convex programs," *Siam Journal of Optimization*, vol. 20, pp. 3427–3464, 2010.
- [34] K. W. Koput, "A chaotic model of innovative search: some answers, many questions," *Organization Science*, vol. 8, no. 5, pp. 528–542, 1997.
- [35] D. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *The Journal of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.

- [36] M. Grant and S. Boyd, “CVX: Matlab software for disciplined convex programming, version 1.21,” <http://cvxr.com/cvx>, Aug. 2010.
- [37] J. M. G. Vilar, H. Y. Kueh, N. Barkai, and S. Leibler, “Mechanisms of noise-resistance in genetic oscillators,” *Proceedings of the National Academy of Sciences of the United States of America*, vol. 99, no. 9, pp. 5988–5992, 2002.
- [38] M. Rathinam, P. W. Sheppard, and M. Khammash, “Efficient computation of parameter sensitivities of discrete stochastic chemical reaction networks,” *The Journal of Chemical Physics*, vol. 132, pp. 034 103–1–034 103–13, 2010.
- [39] J. Löfberg, “Yalmip : A toolbox for modeling and optimization in MATLAB,” in *Proceedings of the CACSD Conference*, Taipei, Taiwan, 2004. [Online]. Available: <http://users.isy.liu.se/johanl/yalmip>